

ASGC:

-----

- \* Complete CRONUS server setup at ASGC: [cronus.grid.sinica.edu.tw](http://cronus.grid.sinica.edu.tw)
- \* Accounting fixed for T1 with help from Dave Kant RAL
- \* Local T2 services migrated to blade servers

#### CERN Tier-0 Site Report

-----

- The Alice Castor Stager database was successfully moved to the DES RAC setup on Wednesday morning. The associated DLF database was also be moved to new hardware.
- A similar intervention for the LHCb Castor databases is being prepared for Monday May 7.
- The dedicated Atlas Tier-0 Castor instance, running new Castor s/w version 2.1.3, is now being used by Atlas. Initial feedback is positive.
- The replacement diskserver hardware to be used in the NA48 CDR has been handed over to the experiment. The Compass setup is progressing slowly.
- Problems on diskserver hardware (both old and new servers) are keeping many people very busy...
- castorlhcb stager had a meltdown due to changed usage of the "durable" lhcbdata pool Wednesday evening. The problem was solved yesterday evening by increasing the number of slots on the diskserver but this probably need further tuning.
- ce103 and ce104 have been drained and are being retired (batch hw).

#### FNAL Tier-1 Site Report (7 May 07)

-----

We are fully operational and much work is being done. The only thing not working are the OPS SAM tests. We've worked on fixing it for more than a day, and have not been able to fix it. The cp seg faults.

Tier1 GridKA (FZK): [ author : Jos van Wezel]

-----

Nothing to report. Trouble free operation.

#### PIC Tier-1 Availability report:

-----

- lcg-CE: The service has been running ok with no major issues.
- SRM-disk and SRM-tape: The service has been running ok with no major issues except some SAM false positive showing SRM problem at the 28 of may from 08:00-10:00 GMT on both SRM [castorsrm.pic.es](http://castorsrm.pic.es) and [srm-disk.pic.es](http://srm-disk.pic.es) due to a SAM problem: BDII Connection Timeout: [sam-bdii.cern.ch:2170](http://sam-bdii.cern.ch:2170).

-glite-CE: ce05.pic.es was not working due to the middleware error "The job attribute PeriodicHold expression ""Matched != TRUE && CurrentTime > QDate + 900"" evaluated to TRUE.

Problem seems to be fixed now by the latest gLite update.

#### IN2P3

-----

1) Too much timeout with the new Top Regional BDII :

A lot of CE-sft-lcg-rm failures was due to Top BDII timeout. After having monitored the requests made to our regional Top BDII, indexes were added to the LDAP databases. The problem seems to be fixed now.

2) CEs unavailability during 2 hours due to an update of the LRMS client. Problem was quickly identified and CEs' jobmanager was consequently modified.

#### NIKHEF

-----

There were no major issues this week.

#### NDGF

-----

Planned network outage on the core servers for the tier1 storage from 20070503 07:00 UTC to 20070503 08:22 UTC

#### RAL

-----

3 LHCb tapes were corrupted due to a hardware reset, this was causing restore requests for those file to hang and the retry of these requests was slowing other requests down. LHCb have been contacted about affected files, however the file s may have been recovered

#### BNL

-----

Friday:

One of seven gridftp door nodes (dcdoor02.usatlas.bnl.gov) was off-line around midnight. The server crashed with an out-of-memory error. The server will be rebooted at 9:30AM.

The maximum TCP buffers for read and write were 4Mbyte. Each file has 10 TCP streams. Therefore the total network buffer allocated to a single file may be up to 40 Mbyte. We will decrease the maximum TCP buffer on one host to prevent the memory problem. We will monitor the host to see whether the data transfer performance will be affected by a smaller buffer size. If the change does not affect the performance, we will adjust all the remaining GridFtp server nodes.

Monday:

Tuesday:

Dantong & Jason were looking into why one development server was powered off and trying to turn it back on and may have accidentally nudged the power cable on the SRM database, which is directly above this development server.

The SRM database was offline at 18:20PM. The RACF operator discovered problems from the Nagios monitoring page, and notified the on-call person. The SRM server was resumed on 9:00PM.

Further investigations show that the electric cable is indeed clipped onto the server. The poor design of cable brackets always wiggles the cables when brackets are opened and closed.

Wednesday:

A new ATLAS switch (CISCO 6513) was added into the USATLAS computing facility. It provides the network connectivity for 120 newly purchase ATLAS worker nodes and other infrastructure servers.

Thursday:

The new ATLAS Switch (sw22) was put into production.

BNL demonstrated a great performance in ATLAS Data Export. There is 300MB/second continuously for several hours.

LCG 3D: BNL 3D instance is ready for production. The plan is to add it into the ATLAS stream production after CERN 3D team fixes the on-going CERN network problem which has brought down the stream replications from CERN to all other Tier 1 sites. No firm date was decided in the 3D weekly meeting due to the network problem.

Friday:

BNL dCache continues providing stable and high performance. The data transfer into BNL dCache has sustained at the speed of 300MB/second for more than 24 hours. The maximum data rate into dCache is about 400M Byte/second for one hour. It consists of both Tier 0 data export exercise and USATLAS production data streams. BNL dCache has demonstrated its best stability and performance.